

Time-dependent evolution of RF-generated nonthermal particle distributions in fusion plasmas

J. C. Wright¹, A. Bader¹, E.D’Azevedo², L.A. Berry², R.W. Harvey³,
E.F. Jaeger⁴, J.-P. Lee¹, C.K. Phillips⁵, A. Schmidt¹, Ky-Anh Tran,¹,
E. Valeo⁵, P.T. Bonoli¹

¹Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA 02139

²Oak Ridge National Laboratory, Oak Ridge, TN 37831

³CompX, Del Mar, CA 92014

⁴XCEL Engineering, Oak Ridge, TN 37830

⁵Princeton Plasma Physics Laboratory, Princeton, NJ 08543

E-mail: jcwright@mit.edu

Abstract. We describe fully self-consistent time-dependent simulations of radio frequency (RF)-generated ion distributions in the ion cyclotron range of frequencies and RF-generated electron distributions in the lower hybrid range of frequencies using combined Fokker-Planck and full-wave electromagnetic field solvers. In each regime, the non-thermal particle distributions have been used in synthetic diagnostic codes to compare with diagnostic measurements from experiment, thus providing validation of the simulation capability. The computationally intensive simulations require multiple full-wave code runs that iterate with a Fokker-Planck code. We discuss advanced algorithms that have been implemented to accelerate both the massively parallel full-wave simulations as well as the iteration with the distribution code. A vector extrapolation method that permits Jacobian free acceleration of the traditional fixed-point iteration technique is used to reduce the number of iterations needed between the distribution and wave codes to converge to self-consistency. The computational burden of the parallel full-wave codes has been reduced by using a more efficient three-dimensional parallel decomposition that improves the strong scaling of the codes and reduces the communication overhead.

1. Introduction

Radio frequency (RF) power in plasmas injected from antennas and waveguides can create flows, modify the magnetic topology with driven currents, and help raise the temperature to trigger fusion reactions through heating of ions and electrons [1]. The amplitudes of externally driven radio frequency (RF) waves for use in fusion-relevant plasmas induce fluctuations in the plasma that are small compared with the background density and fields. In addition, the wave periods are much shorter than the time scales on which the target plasma evolves. The physics of these waves are described by Maxwell’s equations in a dielectric media with the qualification that, because charges are free to move in the plasma, there is no distinction between conduction and polarization currents. The plasma dielectric tensor is formally a nonlocal integral operator derived from the Vlasov equation for the plasma particle distribution as perturbed by the RF fields [2].

The two wave codes discussed in this article, **AORSA** [3] and **TORLH** [4], are applied to the solution of the wave equation (WE) as shown in Eq. (1) for the ion cyclotron range of frequencies (ICRF) and lower hybrid range of frequencies (LHRF), respectively. Formulated in the frequency domain, the plasma RF wave equation is an elliptic system for the rapidly oscillating wave electric field \mathbf{E} . The plasma response is represented by \mathbf{J}^P , which is in general an integral operator linear in the wave electric field. The wave frequency is ω , and \mathbf{J}^P is the plasma current response to the RF. The decomposition used in **TORLH** is given by Eq. (1b), which uses a toroidal coordinate system with a spectral basis in the periodic dimensions θ and ϕ and a finite-element basis for the radial dimension. The **AORSA** code decomposes the fields in spectral basis in all three of coordinates of the cylindrical system (R, Z, ϕ) as shown in Eq. (1c).

$$\nabla \times \nabla \times \mathbf{E} = \frac{\omega^2}{c^2} \left\{ \mathbf{E} + \frac{4\pi i}{\omega} (\mathbf{J}^P + \mathbf{J}^A) \right\} \quad (1a)$$

$$\mathbf{E}(r, \theta, \phi) = \sum_m \mathbf{E}_m(r) \exp(im\theta + in\phi) \quad (1b)$$

$$\mathbf{E}(R, Z, \phi) = \sum_{n,m,l} \mathbf{E}_{n,m,l} \exp(i(k_n R + k_m Z + l\phi)) \quad (1c)$$

$$\mathbf{J}^P(\mathbf{x}) = \sum_i \overset{\leftrightarrow}{\sigma}_c(\mathbf{k}^{(i)}, \mathbf{x}) \cdot \mathbf{E}_i(\mathbf{x}) \quad (1d)$$

The dielectric tensor, $\overset{\leftrightarrow}{\sigma}_c$, is an integral function of the particle distribution function, $f_0(\mathbf{v})$. This distribution function may be specified as a Maxwellian, in which case the dielectric tensor is given in terms of Fried-Conte functions [2, 5] and modified Bessel functions. But for self-consistency, $f_0(\mathbf{v})$ must be evolved in the presence of the RF waves. This is done with the Fokker-Planck(FP) code, **CQL3D** [6] that solves the long time scale response of the distribution function to the RF waves [1] from a bounce-averaged, zero-orbit width Fokker-Planck. A simple fixed-point iteration scheme between the FP and WE codes scripted in python has been used to achieve self-consistency. This technique results in an explicit time advance requiring many expensive evaluations of the WE and small time steps.

In this paper, we discuss the application of a reduced-rank extrapolation (RRE) and minimal polynomial expansion (MPE) method to the iteration technique. We show that in the LHRF where the steady state is of interest, accuracy, and stability can be improved in few iterations with this method. In case of ICRF waves, an implicit time advance can be achieved resulting in significant time savings. In both frequency regimes we discuss validation efforts with experiments through the use of synthetic diagnostics. These comparisons indicate the full-wave LH simulations capture important effects of interference of the waves present in cases where the absorption of these waves by the plasma is relatively weak. In the simulations of ICRF heating, comparisons of the energy spectrum of heated ions motivates the inclusion of more detailed effects of the ion orbits on the evolution of $f_0(\mathbf{v})$ neglected by the FP code that improve agreement with experiment.

2. Physics of LH and IC waves

2.1. Lower hybrid slow waves for current drive

RF electric fields in the lower hybrid range of frequencies are characterized by ($\omega_{\text{LH}} \sim \sqrt{\Omega_{\text{ce}}\Omega_{\text{ci}}}$), where $\Omega_{i,e} \equiv qB/m_{i,e}c$ are the electron and ion cyclotron gyration frequencies in a magnetic field of strength B . This frequency is on the order of several gigahertz magnetic field strengths of several tesla. These waves have perpendicular wavelengths of the order of 1 mm at densities of the order of 10^{20}m^{-3} in a fusion device with scale lengths on the order of 1 m. They are nearly

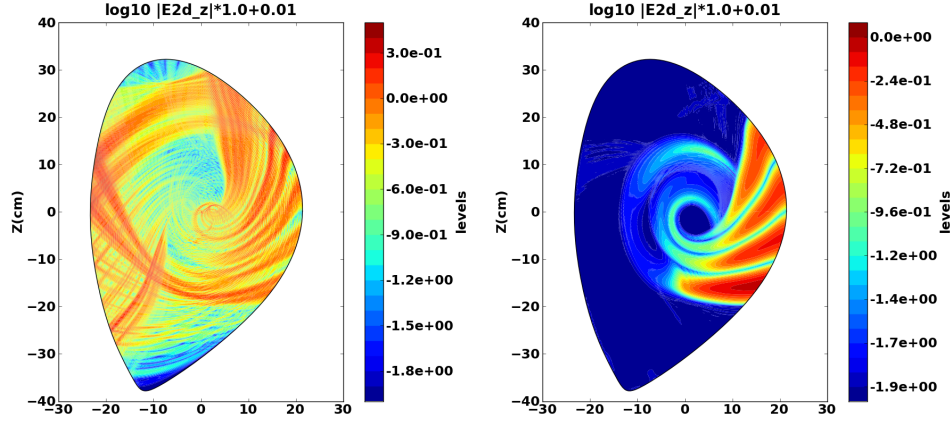


Figure 1. Left panel: TORLH / CQL3D simulation of weak damping regime in Alcator C-Mod ($n_{\parallel} = -1.6$). Right panel: TORLH / CQL3D simulation of strong damping regime in Alcator C-Mod ($n_{\parallel} = -3.1$). The RF electric field parallel to the magnetic field is shown for a single toroidal mode. The electric field units are arbitrary and are driven by a unit amplitude from four waveguides on the middle right of the plasma cross-section. The plasma parameters [$B_0=5.36$ T, $n_{e0}=7 \times 10^{19} \text{m}^{-3}$, $T_{e0}=2.33$ keV] were held constant for the two simulations.

electrostatic with a high-phase velocity and interact with electrons at many times the electron thermal speed, which is given by $v_{te} = \sqrt{2T_e/m_e}$ where m_e and T_e are the electron mass and electron distribution temperatures, respectively, and v_{te} is the electron thermal velocity. Because collisional drag decreases with velocity as $1/v^3$ in plasmas, LH waves are ideal for driving current to control the magnetic configuration. Typically, currents are maintained in fusion plasmas with an external transformer coil with the plasma being the secondary winding, but this is inherently not a steady-state solution.

Because the interaction with electrons takes place at velocities far above their thermal velocity, the distribution function is changed from an equilibrium Maxwellian and is instead determined by the balance between collisional and RF wave-induced diffusion in velocity space. The evolution of the distribution function in turn affects the dielectric of the plasma and increases the damping of the waves. An iterative scheme, described in detail in Section 3.2, is used to achieve self-consistency. In Fig. 1 we present two cases distinguished by the parallel index of refraction used. Here the parallel refractive index (n_{\parallel}) is defined as $n_{\parallel} = c/v_{\parallel}$. Both cases have been iterated to steady state with the CQL3D Fokker-Plank code using fixed-point iteration requiring 30 steps for $n_{\parallel} = -1.6$ and 15 steps for $n_{\parallel} = -3.1$. The damping is stronger for the higher index of refraction (right panel of Fig. 1) corresponding to a lower wave phase velocity, and thus the distribution function doesn't evolve far from equilibrium. The damping is weaker for the lower index of refraction (left panel of Fig. 1) and corresponds to an interaction velocity that is significantly relativistic. In this case we can see the unit driving amplitude is increased approximately a factor of $3 = (10^{0.3})$ in the interior fields because of cavity amplification. Interference structures and reflections are also evident.

2.2. Ion cyclotron fast waves for minority heating

Ion cyclotron resonance heating (ICRH) of the hydrogen-minority species in a deuterium majority plasma is the prime auxiliary heating scheme on Alcator C-Mod. This heating method is also used on many other tokamaks and will be used on ITER. Rather than increasing the parallel energy of electrons as in LH, ICRH increases the perpendicular energy of ions at their

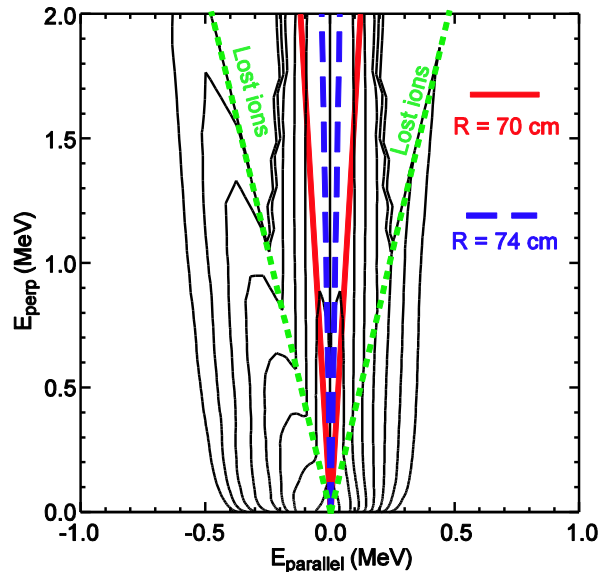


Figure 2. The distribution function resulting from self-consistent coupling between AORSA and CQL3D for a minority heating case shown at a normalized minor radius of $r/a=0.3$. The green dashed lines indicate the region in momentum space where ions become trapped. The solid red and dash blue lines show which trapped ions are viewed by the CNPA ion diagnostic discussed in Section 4.2.

cyclotron resonance. The effect is shown in Fig. 2, where the contours of the ion distribution function are strongly distorted along the vertical axis associated with perpendicular energy. The characteristic “rabbit ear” structure is a consequence of the poloidal variation of the magnetic field that creates a magnetic bottle in that direction with consequentially trapped and lost ions determined by their perpendicular energy and magnetic moment. Heating the minority species with an ICRF wave at the minority ion species fundamental cyclotron frequency is more efficient than heating the majority species because the plasma polarizes the ICRF wave opposite of the ion gyration at the location of the majority cyclotron resonance [7]. Energetic ions already present in the plasma or generated by the ICRH itself can influence wave absorption by increasing single pass absorption, modifying the deposition profile and changing the energy partition between electrons and ions. We simulate ICRH interaction with the plasma by coupling the spectral full-wave solver AORSA with the zero ion-orbit-width Fokker-Planck code CQL3D. Coupling the codes allows self-consistent solutions that include the effects of energetic ions on the fields [8].

3. Numerical Techniques

3.1. Higher-order parallelism

We have developed the tools for solving for the self-consistent RF electric fields in the lower hybrid frequency regime [9, 10]. The full-wave code TORLH, a modified version of the TORIC ICRF code [11], has been made more efficient with a new parallel solver [4] that reduces the execution time a factor of 3–5 for typical runs on 64 to 128 processor cores on the order of one to two hours. The stiffness matrix that results from the solution of the system in Eq. (1) using the decomposition in Eq. (1b) results in a block tridiagonal system. Previously, the parallel decomposition was done for each block in the system in turn; that is, all processors were used in the inverse of each block. As shown in Fig. 3, this new solver scales to tens of thousands of cores efficiently enabling wall clock times of several minutes. The improvement is due to an avoidance of the strong scaling limit of the direct solve of the blocks by extending the parallelism of the solver along the radial dimension (associated with block rows of the stiffness matrix). This is done by using a combination of cyclic reduction and divide and conquer algorithms that are common in tridiagonal solvers [12]. A recent upgrade of the CQL3D code [6] now employs an implicit time advance and is parallel in the radial dimension, bringing run times of this Fokker-Planck code down to the order of a few minutes on 32 cores for problems discussed in this paper.

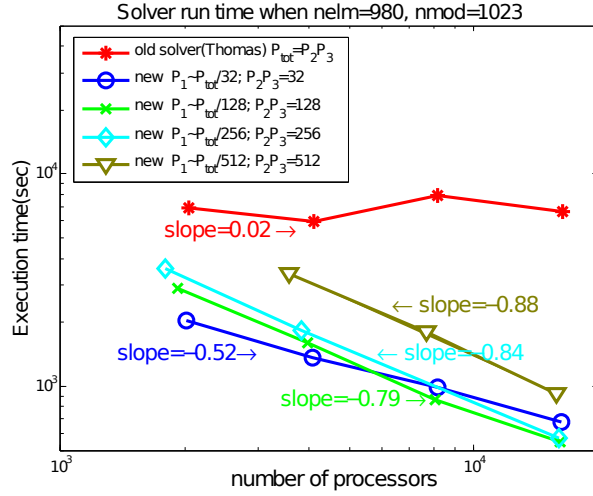


Figure 3. Comparison of execution time of parallel solver for different configurations of the three dimensional processor mesh. P_1 is the number of processors in the block row direction. P_2P_3 is the number of processors in the two dimensional decomposition of the blocks. The four decompositions shown exhibit some variation in execution time. The three dimensional decomposition is in general 5–10 times faster than the Thomas solver and demonstrates much better parallel scaling.

Iteration of these codes produces an electron distribution function and a wave electric field that are consistent with one another.

One way to implement a 3-D grid is to use a context array in BLACS (Basic Linear Algebra Communication Subprograms) in which each context uses a 2-D processor grid as the 2-D solver. In BLACS, a context indicates a group within a boundary of an MPI communicator. In the default context having the total number of processors, one can assign several subgroups of processors corresponding to each context according to the specific maps. The subgroups can communicate with each other when needed in a tridiagonal algorithm. BLACS provides a convenient programming interface for the decomposition, but it actually uses sub-MPI COMMUNICATORS to achieve this, and therefore the same algorithm could be implemented in pure MPI if desired. But BLACS is a requirement for ScaLAPACK [13] and so use of the library here adds no new dependencies.

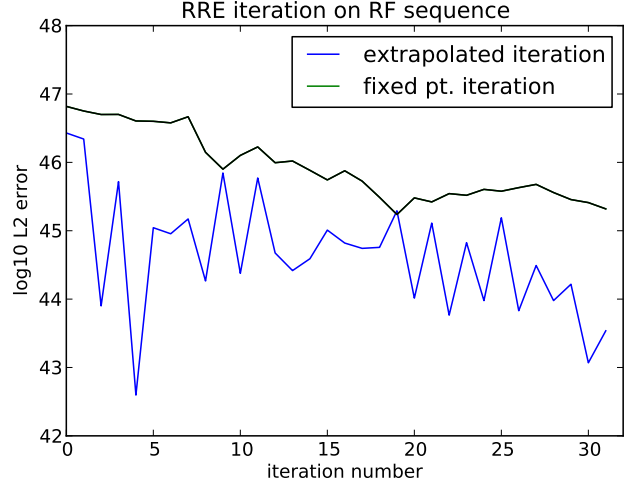
3.2. Vector acceleration

Equation 2 describes the normal fixed-point iteration technique used to iterate between a wave equation (WE), which determines the wave electric field, and the related quasilinear diffusion in velocity space (D_{ql}) as described by Eqs. (1) and Fokker-Planck (FP) equation, which advances the ion or electron particle distribution (f) subject to collisional and wave induced diffusion (D_{ql}). The values at the $(n + 1)$ th iteration are given simply from the operators FP and WE at the n th step:

$$\begin{aligned} f_{n+1} &= FP(f_n, D_{ql_n}) \\ D_{ql_{n+1}} &= WE(f_n). \end{aligned} \quad (2)$$

This technique is known to converge slowly in many instances. It can also be unstable in cases of weak ($n_{\parallel}^2 < 40/T_e$) absorption unless used carefully (small time steps in the Fokker-Planck code and ramping up of power slowly.) A vector extrapolation method [14, 15] that permits Jacobian-free acceleration of the traditional fixed-point iteration technique is used to reduce the number of iterations needed between the distribution and wave codes to converge to self-consistency. This method requires only the solutions resulting from the iterations of the fixed point method and refines the solution to much higher accuracy at negligible cost, as shown in Figure 4. The solution from fixed-point iteration must be sufficiently close to the correct solution to linearize about the exact solution to first order. The unknown Jacobian is then approximated through a

Figure 4. Acceleration of high-phase velocity LH iteration implemented in python. The L2 norm of the difference of Dql between successive iteration steps is plotted on a \log_{10} scale. Results have not been normalized to vector length. Refinement of solution takes less than 10 minutes, with a gain in accuracy of more than 50 times ($10^{1.7}$).



least-squares inversion using the iteration vectors generated from the fixed-point method. This determines the weights to be used to reconstruct a more accurate approximation from a linear superposition of the iteration vectors.

3.3. Implicit coupling

In the previous section vector acceleration with the RRE technique was used to refine the solution of a fixed-point iteration. In this section we apply the minimum polynomial expansion (MPE) version of vector acceleration to extend an explicit time advance of a nonlinear system to an implicit one. The FP-RF couplings are nonlinear systems with both integral and differential terms. Explicit coupling often requires undesirably small time steps. Using the **AORSA** and **CQL3D** codes in the ICRF regime, we use MPE to resolve the evolution of the ion distribution function in time with an implicit time advance. This is done without modifying or using knowledge of the internal algorithms of the two codes and has been implemented in the Integrated Plasma Simulator (IPS) of the CSWIM SciDAC Project [16].

$$\mathbf{E}_{i+1} = \text{AORSA} \left(\text{CQL3D}_{\Delta t} \left(f_i, \frac{D_{QL}^i(\mathbf{E}_i) + D_{QL}^{i+1}(\mathbf{E}_{i+1})}{2} \right) \right) \quad (3)$$

Conceptually, the implicit solve is a centered Crank-Nicholson advance as given in Eq. (3). The centering of the quasilinear operator is accomplished by using the MPE algorithm to correct the fixed point iteration between **AORSA** and **CQL3D** to get D_{QL}^{i+1} . Then $D_{QL}^{i+1/2}$ is used to make the implicit advance of the electric field. The entire process is illustrated in Fig. 5. In Fig. 6, the calculation of the hydrogen distribution energy for the explicit and implicit time advance are compared for different time steps. It is seen that the implicit advance with 6 time steps reproduces the curve using 160 explicit time steps. The MPE implicit advance required 2981 CPU-hours (2.3 hours on 1,296 CPUs) and the standard fixed point explicit advance used 10368 CPU-hours (8 hours on 1,296 processors.)

4. Synthetic Diagnostics and Model Validation

4.1. Hard x-ray emissions for LHRF

On Alcator C-Mod, a hard x-ray (HXR) camera diagnostic is used to indirectly measure the energy of LHRF-accelerated electrons from the bremsstrahlung radiation emitted (via the “headlight” effect) as they circle the toroidal device—much like in a circular particle accelerator.

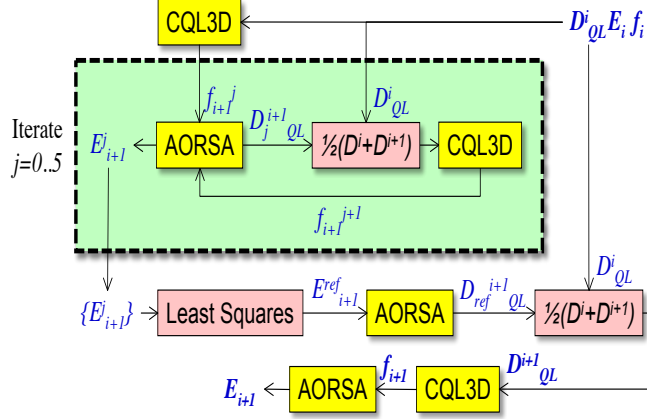


Figure 5. Flow chart showing the implicit advance of the distribution function and electric fields using a centered RF diffusion operator. Steps in the shaded region use a small number of explicit advances (j index) in conjunction with vector acceleration to produce a time centered RF diffusion (D_{ql}) that is used in the outer implicit advance.

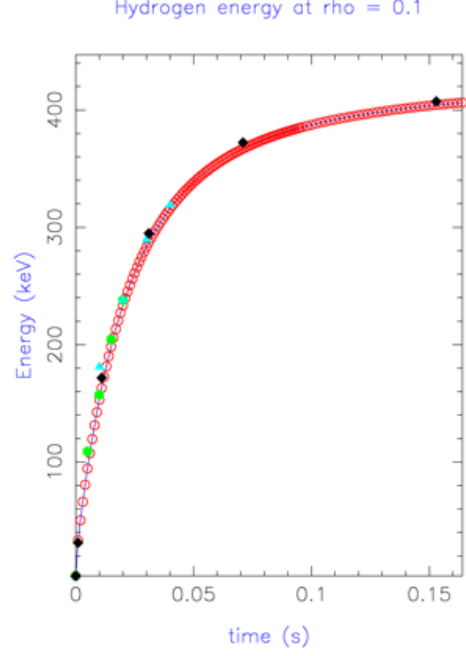


Figure 6. Implicit versus explicit time advance for evolution of minority ion distribution using CQL3D-AORSA codes. The maximum energy on the mesh was 2000 keV. The toroidal mode was $n_\phi = 10$, and the RF power was 2.65 MW.

When the fast electrons pitch angle scatter off of back ground ions and electrons their conical beam of X-rays can be directed along the viewing chords. The diagnostic consists of 32 chords viewing horizontally through a poloidal plane in the plasma and focusing to a detector outside of the plasma [17]. This diagnostic is simulated in the CQL3D code from the calculated electron distribution and is used as a basis of comparison with the experiment.

When the absorption strength of LH waves is weak, the waves are in a multipass regime where interference of crossing waves, reflections from the cutoff and caustic formation, must be treated properly. In this case, it is reasonable to expect significant differences between full-wave and ray-tracing predictions. In Figure 7, we compare the HXR predictions for multipass regime using self-consistent TORLH /CQL3D simulations for the evolution of $f_e(\mathbf{v})$ and calculation of the HXR signal. The HXR signal from ray tracing in Figure 7 is a factor of 5 larger than the experiment [18], in contrast to the HXR signal from the full-wave analysis. This may be due to the differences in the reconstructed Dql from the two approaches, since the Dql from full wave includes phase interference effects that can lead to a reduction in the energy density that are absent in first-order WKB ray tracing.

4.2. Compact neutral particle analyzer for ICRF

The fast-ion distribution created during minority ion cyclotron heating is measured on the Alcator C-Mod device with a compact neutral particle analyzer (CNPA) [19]. The CNPA consists of vertically viewing photodiode detectors with viewing chords from $r/a = 0.18$ to 0.60 and energies from 150 keV to 1.5 MeV. The CNPA measurements are compared with the

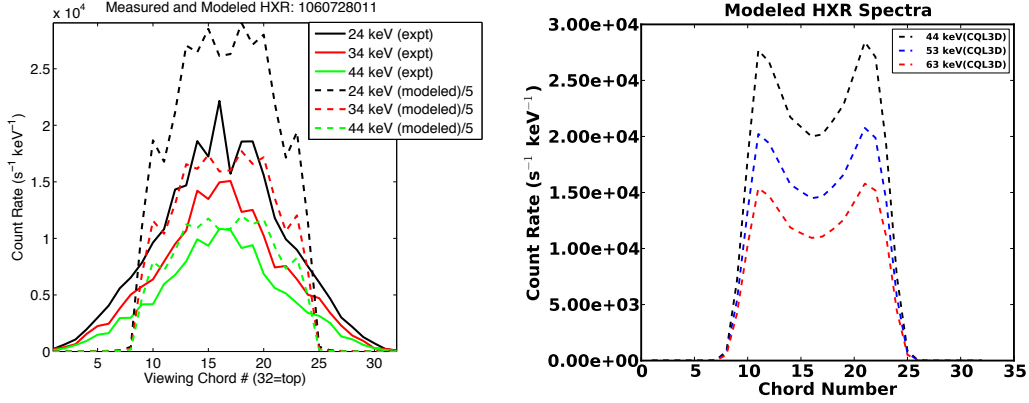


Figure 7. Comparison of measured and modeled HXR profiles for shot 1060728011. Modeled profiles from ray tracing (left panel) have been scaled by a factor of 1/5 and those from full wave (right panel) have not been scaled. The HXR camera views horizontally with chord 1 (32) corresponding to the bottom (top) of the tokamak. Full-wave simulations assume 500kW coupled LH power. The plasma parameters are $[B_0=5.36 \text{ T}, n_{e0}=7 \times 10^{19} \text{ m}^{-3}, T_{e0}=2.33 \text{ keV}]$.

output of a synthetic diagnostic that has been added to the Fokker-Planck solver CQL3D [6]. The synthetic diagnostic calculates the flux of neutrals (in events per second per steradian per cm^2 per keV) along a sightline. Fast ions can neutralize by charge exchange with background neutral deuterium or hydrogen-like boron or by recombination with thermal electrons. Boron is a common impurity in C-Mod, and charge exchange with boron is the dominant neutralization process for most of the energy range that we measure. The densities of the D^0 and the B^{4+} are specified as inputs to the synthetic diagnostic. These are not measured directly but are instead calculated by the transport analysis code TRANSP (for D^0) and by a collisional equilibrium calculation (for B^{4+}), where the B^{5+} density is assumed to be a constant percentage of the electron density. This percentage is unknown, although spectroscopic measurements estimate it to be $\sim 1\%$ of the electron density, n_e . In order to produce agreement with the simulation, boron densities of 2.5% are used for all simulations presented in this paper. The synthetic diagnostic is found to be insensitive to variations in electron density, electron temperature or hydrogen fraction, provided that hydrogen fraction is low enough that the waves are not in the strongly mode-converted regime. However, the synthetic diagnostic is very sensitive to sightline placement near the resonance layer. This sensitivity prompts the use of multiple sightlines in order to model a sightline with finite width.

Simulations were performed by iterating AORSA and CQL3D by fixed point iteration until a self-consistent fast-ion distribution is obtained. Time-dependent simulations were advanced in steps of 1 ms in the CQL3D code with substeps of 0.1 ms. Four simulations of four plasma discharges at currents ranging from 0.6 to 1.2 MA were performed and the results show good agreement between simulation and the CNPA experimental results for all currents. Shown in the left panel of Fig. 8 are results from a simulations at 1.0 MA. The simulations correctly predict the energy dependence of the fast-ion signal and the relative intensity of signal at the two channels. There is excellent agreement in the steady-state case (left panel) in which the simulation correctly reproduces the energy dependence of the fast ion signal. Furthermore, the simulation correctly predicts the relative strength of the signal on the two channels, showing that it captures the better fast-ion confinement in the high-current plasmas [19].

Though good agreement is obtained between simulation and experiment for fast-ion

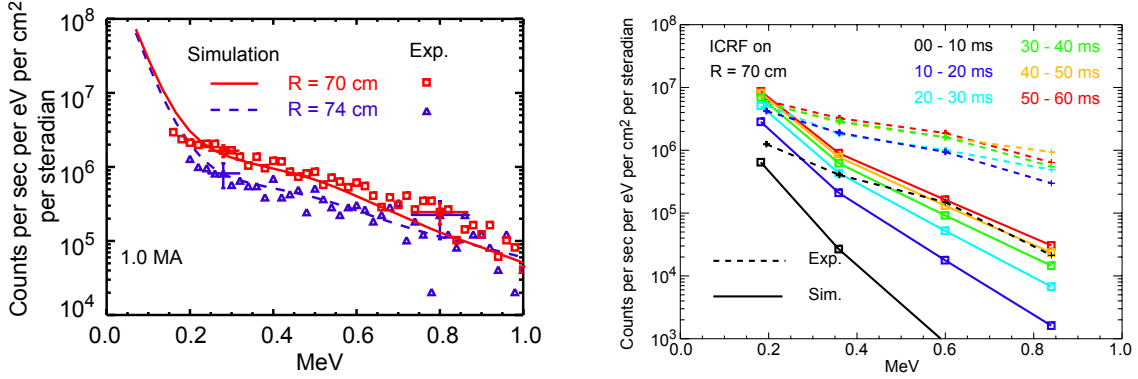


Figure 8. NPA counts: comparisons between simulation and experiments are shown for 1.0 MA in steady state (on the left) and while evolving (on the right.)

distributions that are in steady state ($df/dt = 0$), there are discrepancies between the experimental and simulated results for the time-dependent evolution of the fast-ion distribution. In particular, the experimentally measured rise time of the ion tail is faster than the simulated rise time, which can be seen in the measured spectra shown as dashed lines in the right panel in Fig. 8. Several reasons for this discrepancy were investigated: finer time-stepping algorithms do not improve the simulation results, the faster rise and decay cannot be explained by changes in boron concentration, and increasing radial diffusion can produce a faster rise and decay, but cannot produce the correct shape with respect to energy regardless of what diffusion parameters are chosen. A possible explanation of the discrepancy is due to violation of certain assumptions of quasilinear theory, which assumes successive interactions between the ions and the ICRF waves are uncorrelated and that the “kick” that an ion receives when passing through resonance with the wave produce a small percentage change in its energy.

To address the discrepancy in rise time of different energy bins, we use the particle code DC [20] to calculate the diffusion coefficient directly from single particle motion. Particles are launched equispaced in initial gyro phase about a given gyro center and equispaced in toroidal length. The diffusion coefficients are obtained by averaging the resulting square of the velocity changes after one (or more) particle poloidal circuits, to obtain the ICRF bounce-averaged diffusion tensor. This averaging is carried out for a 3D array ($u_{\parallel}, u_{\perp}, R$) of initial conditions, giving the six independent RF diffusion coefficients in 3D constant-of-motion space. The method follows the formalism of Refs. [21, 22]. To isolate the differences with quasilinear theory, the particle orbits were adjusted with fictitious forces that compensated for drifts of particle motion due to magnetic field gradients and curvature. This provides a fair basis of comparison with the zero orbit width model used in calculating D_{ql} from the wave code. The integration of 268M Lorentz orbits ($32 \text{ radii} \times 128 \text{ } u_{\perp} \times 256 \text{ } u_{\parallel} \times 8 \text{ gyro phase} \times 32 \text{ toroidal angle starting positions}$) is trivially parallel and takes 7 hours on 4096 cores.

Recall that the right panel of Fig. 8 are results from the time-resolved evolution of the ion distribution at three time points during the rise of ion energy as shown in Fig. 6. The vertical perpendicular viewing chords on Alcator C-Mod are particularly sensitive to the minority-ion perpendicular velocity distribution, which is most affected by the enhanced pitch angle scattering calculated with DC due to the retention of correlation effects. Comparing Fig. 9 with the right panel of Fig. 8, we see that the DC results bring the energy dependence of the evolution of the distribution function into better agreement with experimental measurements.

Figure 9. NPA energy spectra from CQL3D using AORSA QL coefficients (left) and using DC Lorentz orbit based coefficients, at $t=4$ msec. Results are for three vertical view cords and one horizontal chord (lowest, at low energy). The DC coefficient case has a relatively high temperature tail and less difference between viewing chords, as a result of the more filled-in perpendicular distribution.

5. Conclusions

We have presented simulation results from a combined LH full-wave/2-D (u_{\parallel}, u_{\perp}) Fokker-Planck calculation for parameters characteristic of the weak damping regime in the Alcator C-Mod device. Accurate treatment of the evolution of the distribution function consistent with the wave fields was necessary to achieve these results. We employed advanced numerical algorithms to accelerate the convergence of the required iteration procedure and showed that the computational burden can be reduced a factor of 4 to 5 and/or the accuracy can be increased a factor of approximately 50. Simulated HXR spectra have been computed from the converged electron distributions predicted by CQL3D and validated against experimentally measured spectra.

AORSA-CQL3D simulations have been performed for Alcator C-Mod minority-heated plasmas and the simulations have been validated by comparing the predictions of a neutral particle analyzer synthetic diagnostic with experimental measurements. The results were found to be in good agreement for steady-state plasmas. However, more rigorous time-dependent simulations were found to evolve much slower than the fast ions in the experiment, indicating that quasilinear theory may be violated. Large parallel calculations of the diffusion coefficient directly from the particle orbits result in an evolution of the ion energies that is closer in agreement with the experiment.

For two frequency ranges of radio frequency waves used in fusion devices, we have done integrated modeling simulations employing advanced mathematical and computational algorithms to introduce new levels of physics fidelity in these models. Through the use of synthetic diagnostics to provide a basis of direct comparison with experiment, we have validated these approaches and shown that they improve their predictive capability.

Acknowledgments

Funding for this research was provided in part under USDOE Contract No. DE-FC02-01ER54648 and the Massachusetts Institute of Technology Undergraduate Research Opportunities Program.

References

- [1] Bonoli P T, Batchelor D B, Berry L A, Choi M, D'Ippolito D A, Harvey R W, Jaeger E F, Myra J R, Phillips C K, Smithe D N, Tang V, Valeo E, Wright J C, Brambilla M, Bilato R, Lancellotti V and Maggiora R 2007 *J. Phys.: Conf. Ser.* **78** 012006 URL <http://stacks.iop.org/1742-6596/78/012006>
- [2] Stix T H 1992 *Waves in Plasmas* chap 10, p 250 in stixwaves [7]
- [3] Jaeger E F, Berry L A, D'Azevedo E, Batchelor D B, Carter M D and White K F 2002 *Phys. Plasmas* **9** 1873–1881
- [4] Wright J C, Valeo E J, Phillips C K, Bonoli P T and Brambilla M 2008 *Commun. Comput. Phys.* **4** 545–555
- [5] Fried B and Conte S 1961 *The Plasma Dispersion Function* (New York: Academic Press, New York)
- [6] Harvey R W and McCoy M G 1993 *Proc. of the IAEA Tech. Committee Meeting (Montreal, 1992)* IAEA (Vienna: Institute of Physics Publishing; USDOC/NTIS Doc. DE93002962) pp 489–526
- [7] Stix T H 1992 *Waves in Plasmas* (New York: American Institute of Physics)
- [8] Jaeger E F, Berry L A, Ahern S D, Barret R F, Carter M D, D'Azevedo E, Moore R D, Harvey R W, J R Myra D A D, Dumont R J, Phillips C K, Smithe H O D N, Bonoli P T, , Wright J C and Choi M 2006 *Phys. Plasmas* **13** 056101–056107
- [9] Wright J C, Bonoli P T, Schmidt A E, Phillips C K, Valeo E J, Harvey R W and Brambilla M A 2009 *Phys. Plasmas* **16** 072502
- [10] Valeo E J, Phillips C K, Bonoli P T and Wright J C 2007 *17th Topical Conference on Radio Frequency Power in Plasmas* 933 ed Ryan P and Rasmussen D (New York: American Institute of Physics) p 297
- [11] Brambilla M 1999 *Plasma Phys. Controlled Fusion* **41** 1–34
- [12] Garaud P and Garaud J 2008 *Mon. Not. R. Astron. Soc.* **391** 1239–1258
- [13] Blackford L S, Choi J, Cleary A, D'Azevedo E, Demmel J, Dhillon I, Hammarling S, Henry G, Petitet A, Stanley K, Walker D and Whaley R C 1997 *ScaLAPACK user's guide* (Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, Editor: Dongarra, Jack J.) ISBN 0-89871-397-8
- [14] Sidi A 1991 *J. Comput. Appl. Math.* **36** 305 – 337 ISSN 0377-0427
URL <http://www.sciencedirect.com/science/article/pii/037704279190013A>
- [15] Sidi A 2008 *Comput. Math. Appl.* **56**(1) 1–24 ISSN 0898-1221
URL <http://portal.acm.org/citation.cfm?id=1371266.1371471>
- [16] Elwasif W R, Bernholdt D E, Berry L A and Batchelor D B 2007 *HPC-GECO/CompFrame – Joint Workshop on HPC Grid Programming Environments and Components and Component and Framework Technology in High-Performance and Scientific Computing* (Montreal, Canada)
- [17] Liptac J, Parker R, Tang V, Peysson Y, and Decker J 2006 *Rev. Sci. Instrum.* **77** 103504
- [18] Schmidt A 2011 *Measurements and Modeling of Lower Hybrid Driven Fast Electrons on Alcator C-Mod* Ph.D. thesis Massachusetts Institute of Technology
- [19] Bader A, Bonoli P, Granetz R, Harvey R, Jaeger E, Parker R and Wukitch S 2011 *RADIO FREQUENCY POWER IN PLASMAS: 19th Topical Conference on Radio Frequency Power in Plasmas* (AIP Conference Proceedings no to be published) ed Phillips C and Wilson R (New York: American Institute of Physics)
- [20] Harvey R, Petrov Y, Jaeger E and Group R S 2011 *RADIO FREQUENCY POWER IN PLASMAS: 19th Topical Conference on Radio Frequency Power in Plasmas* (AIP Conference Proceedings no to be published) ed Phillips C and Wilson R (New York: American Institute of Physics)
- [21] Kaufman A N 1972 *Phys. Fluids* **15** 1063
- [22] Eriksson L and Helander P 1994 *Phys. Plasmas* **1** 308